

METHOD AND APPARATUS FOR MULTI-NODAL, THREE-DIMENSIONAL IMAGING

Cross Reference to Related Application

- 5 [0001] The benefit of the filing date of U.S. application No. 60/273,830 filed on 8 March 2001, which is incorporated herein by reference, is claimed herein.

Technical Field

- 10 [0002] This invention relates to the field of machine vision and, more generally, to determining the positions of objects or features in three-dimensional space. The invention provides methods and apparatus for determining positional information for objects. In particular, the invention relates to determining relationships between local frames of
15 reference for multiple measurement units.

Background

- [0003] A three-dimensional position-sensing device can determine the three-dimensional locations of objects within a limited field of view.
20 One example of such a device is the DIGICLOPS™ stereo vision system available from Point Grey Research of Vancouver, Canada. Such systems are used in a wide variety of applications, such as computer vision systems, tracking applications, object dimensioning applications and others.

25

- [0004] A typical stereo vision camera system comprises two spaced-apart digital cameras. Some prior art stereo vision camera systems have three cameras. Figure 1 shows a two-camera vision system

20080802 12:33 FAX 804 681 4081

- 2 -

1 having two cameras 11A and 11B. The distance b between cameras
11A and 11B is referred to as the "baseline". Each of cameras 11A and
11B has an associated optical axis 16A and 16B and a field of view 12A
and 12B. These fields of view 12A and 12B overlap one another in
5 region 13, which is referred to as the "stereo measurement field".

[0005] The position in a local three-dimensional coordinate system
of a point on an object in the stereo measurement field can be determined
by triangulation. This may be done by locating the point in images from
10 each of cameras 11A and 11B. The position of the point is related to the
(u, v) positions at which the point appears in the two images by a function
as follows:

$$\vec{p}(x, y, z) = F(u_1, v_1, u_2, v_2) \quad (1)$$

where (u_1, v_1) is the position of the image of the point in the image
15 obtained by camera 11A; (u_2, v_2) is the position of the image of the point
in the image obtained by camera 11B and (x, y, z) is the location of the
point in the reference frame of stereo measurement unit 10. The function
 F depends upon the distance b , the optical characteristics of cameras 11A
and 11B, and the spacing between light-sensing elements in cameras 11A
20 and 11B. Stereo measurement unit 10 may be calibrated (i.e. the function
 F can be determined) in any suitable manner including by using any of a
number of prior art techniques known to those skilled in the art. The
coordinate system (x, y, z) is local to vision system 1. If system 1 is
moved relative to an object then the apparent position of the object in
25 coordinate system (x, y, z) will change.

- 3 -

[0006] System 1 includes a processor 14 which receives images from each of cameras 11A and 11B, performs feature extraction to identify corresponding points in each of the images and uses the known function F to determine the locations of the corresponding points in the (x, y, z) coordinate system of stereo measurement unit 10. Since the three-dimensional imaging system 1 employs triangulation techniques, if cameras 11A and 11B remain fixed relative to one another, calibration will be preserved. Movements of cameras 11A and 11B or changes in baseline b can cause spurious measurement results.

[0007] The size of stereo measurement field 13 and the three-dimensional imaging resolution of stereo measurement unit 10, may be improved to some degree by one or more of:

- changing fields of view 12A and 12B of cameras 11A and 11B;
- increasing the resolution of cameras 11A and 11B; and,
- changing the baseline b .

However, each of these techniques has limitations.

[0008] Changing the field of view of cameras 11A and/or 11B may increase the size of stereo measurement field 13 and improve the measurement accuracy of stereo measurement unit 10 for a particular range of distances, but such a change simultaneously decreases the measurement accuracy at other distances. For example, widening the field of view of camera 11A increases the size of stereo measurement field 13 and improves the measurement accuracy of stereo measurement unit 10 for close objects, but decreases accuracy for objects that are farther away.

- 4 -

[0009] Increasing imaging resolution of cameras 11A and 11B improves the three-dimensional imaging resolution of stereo measurement unit 10, but increases the amount of data to be processed and decreases system speed.

[0010] In prior art systems like system 1, there is a trade-off between the size and location of stereo measurement field 13 and the accuracy with which three-dimensional positions can be determined.

10 Increasing baseline *b* (i.e. moving cameras 11A and 11B farther apart) increases the accuracy of three-dimensional measurements made using stereo measurement unit 10. However, increasing the length of baseline *b* causes the location of overlap between the limited fields of view (12A and 12B) of cameras 11A and 11B to move to a region further away from

15 stereo measurement unit 10. This eliminates the ability to determine the locations of closer objects.

[0011] Increasing the length of baseline b may also lead to system calibration difficulties. If cameras 11A and 11B are too far apart, it becomes more difficult to keep cameras 11A and 11B from moving in relation to one another. Increasing baseline b also increases the overall size of stereo measurement unit 10.

[0012] There is a need for improved apparatus and methods for
25 determining the locations of objects.

- 5 -

Summary of the Invention

[0013] The invention relates to methods and apparatus for implementing multi-nodal, three-dimensional imaging systems.

5 [0014] One aspect of the invention provides methods for calibrating three-dimensional imaging systems which include first and second stereo measurement units. The first stereo measurement unit collects information from which the locations of object features within a first stereo measurement field can be determined in a first local coordinate
10 system. The second stereo measurement unit collects information from which the locations of object features within a second stereo measurement field can be determined in a second local coordinate system. In some embodiments, the first and second stereo measurement fields overlap one another in an overlap region.

15 [0015] The calibration methods include methods which determine locations of a plurality of calibration features in both the first and second local coordinate systems. Using the locations of the plurality of calibration features, the methods derive a set of one or more
20 transformations between the local coordinate systems and a common coordinate system. The set of one or more transformations may be used to express locations of object features found in either of the stereo measurement fields in the common coordinate system.

25 [0016] The common coordinate system may be the same as one of the local coordinate systems.

20090805 12:36 FAX 804 681 4081

- 6 -

[0017] In some embodiments, the first and second measurement fields overlap within an overlap region and the plurality of calibration features is located in the overlap region. In some embodiments, the plurality of calibration features is located in the measurement field of one of the measurement units and has a known spatial relationship to another plurality of calibration features located in the measurement field of another one of the measurement units.

[0018] Other aspects of the invention provide other methods, apparatus and computer program products.

[0019] Further aspects of the invention and features of specific embodiments of the invention are described below.

15 Brief Description of the Drawings

[0020] In drawings which illustrate non-limiting embodiments of the invention,

Figure 1 is a schematic representation of a prior art three-dimensional imaging system;

20 Figure 2 is a schematic representation of a three-dimensional imaging system according to one embodiment of the present invention;

Figure 3 shows the data flow between individual imaging devices, and the processor involved in calibration and data integration;

25 Figure 4 is a flow chart depicting an embodiment of the calibration process for a three-dimensional imaging system;

20090522600F

- 7 -

Figure 5 is a flow chart depicting an embodiment of the feature localization and extraction process for a particular stereo measurement unit according to the present invention;

Figure 6 is a flow chart depicting an alternative embodiment of the calibration process for a three-dimensional imaging system;

Figure 7 depicts a preferred embodiment of a stereo measurement unit;

Figure 8 is a three-dimensional imaging system employing three stereo measurement units;

Figure 9 is a schematic representation of a three-dimensional imaging system according to one embodiment of the invention, where the stereo measurement units are not aligned;

Figure 10 is a schematic representation of a three-dimensional imaging system according to another embodiment of the invention, where the stereo measurement units are not aligned;

Figure 11 is a flow chart depicting an embodiment of the transformation calculation process for a three-dimensional imaging system;

Figure 12 is a schematic representation of an imaging system according to an alternative embodiment of the invention showing the data flow between individual imaging devices, and the processor involved in calibration and data integration;

Figure 13 is a schematic representation of an imaging system comprising a pair of measurement units having non-overlapping measurement fields; and,

Figure 14 is a flowchart depicting a method according to the invention.

20080802 12:37 FAX 804 681 4081

- 8 -

Description

[0021] Throughout the following description, specific details are set forth in order to provide a more thorough understanding of the invention. However, the invention may be practiced without these particulars. In other instances, well known elements have not been shown or described in detail to avoid unnecessarily obscuring the invention. Accordingly, the specification and drawings are to be regarded in an illustrative, rather than a restrictive, sense.

10

[0022] This invention provides position measuring systems, which combine the measurement fields of two or more measurement units to provide a larger combined measurement field. The positions of reference points in the combined measurement field may be expressed relative to a common reference frame. The invention also provides methods for calibrating such systems.

15

[0023] Figure 2 is a schematic representation of the principal components of a multi-nodal, three-dimensional imaging system 20 according to one embodiment of this invention. Three-dimensional imaging system 20 comprises at least two stereo measurement units (50 and 60). Each stereo measurement unit 50 comprises at least two spaced-apart sensors. Each stereo measurement unit is capable of acquiring information from which locations of objects in space can be determined. In the embodiments described below, the sensors comprise imaging devices which may comprise digital cameras.

20

25

203060" 6/22/00T

- 9 -

[0024] Any suitable three-dimensional measurement units could be used as, or in place of, stereo measurement units 50. For example, three-dimensional stereo or non-stereo measurement units which include other types of sensors such as infra-red detectors, acoustic detectors, radar
5 detectors and/or laser range-finders could be used.

[0025] In the embodiment of Figure 2, stereo measurement unit 50 comprises two spaced-apart imaging devices 52A and 52B, which have optical axes 54A and 54B respectively. Imaging device 52A has a field
10 of view covering a three-dimensional region 56A. Imaging device 52B has a field of view covering a three-dimensional region 56B. Regions 56A and 56B overlap within a three-dimensional first stereo measurement field 58.

[0026] Similarly, stereo measurement unit 60 comprises two spaced-apart imaging devices 62A and 62B, with optical axes 64A and 64B. Imaging device 62A has a field of view covering three-dimensional region 66A. Imaging device 62B has a field of view covering three-dimensional region 66B. Regions 66A and 66B overlap within a three-
20 dimensional second stereo measurement field 68.

[0027] Stereo measurement units (50 and 60) are oriented such that first and second stereo measurement fields 58 and 68 overlap in a three-dimensional region 70. Any object within the volume of region 70 is
25 within the stereo measurement fields of both of stereo measurement units 50 and 60. In this description, region 70 is sometimes referred to as an "overlapping stereo measurement field" or an "overlap region".

20090901 15:22:00

- 10 -

[0028] In Figure 2, the optical axes (54A, 54B, 64A and 64B) of imaging devices 52A, 52B, 62A, and 62B are illustrated as being generally parallel to one another. This is not a general requirement. For example, Figure 9 shows a case where stereo measurement units 50 and 60 are oriented so that the optical axes of their imaging devices are at an angle Φ to one another. In Figure 9, Φ is approximately 90 degrees. This is not necessary, however. The locations and orientations of stereo measurement units (50 and 60) may be selected to give a desired shape to overlapping stereo measurement field 70.

[0029] Individual imaging devices 52A and 52B within stereo measurement unit 50 are spaced apart by a baseline b_{50} . Similarly, individual imaging devices 62A and 62B within stereo measurement unit 60 are spaced apart by a baseline b_{60} . In addition there is a baseline B which has a length corresponding to the translational distance between stereo measurement unit 50 and stereo measurement unit 60. In most applications (but not always), the baseline distances b_{50} and b_{60} (within stereo measurement units 50 and 60) are shorter than the length of baseline B (between stereo measurement units 50 and 60).

[0030] Assuming that each of stereo measurement units 50 and 60 is calibrated (i.e. the function F of equation (1), or an equivalent thereof, is known for each of stereo measurement units 50 and 60), then the locations of reference points which are located in the stereo measurement field of one of stereo measurement units 50 and 60 can be determined, relative to a local frame of reference, from data obtained by the imaging

- 11 -

devices of that stereo measurement unit. The reference points may comprise objects, parts of objects, features of objects or the like. The locations of reference points on objects may be used to determine characteristic features of the objects. Such characteristic features of
5 objects may include, without limitation, orientation, size, locations and shapes of edges and the like.

[0031] The calibration methods explained below permit system 20 to determine the positions of reference points and/or characteristic
10 features of objects which are located within a combined measurement field relative to a common frame of reference (i.e. a common coordinate system). The combined stereo measurement field over which system 20 may be operative to determine the positions and/or features of objects comprises the union of stereo measurement fields 58 and 68. The
15 common frame of reference may be the same as a local frame of reference associated with one of stereo measurement systems 50 or 60 or may be a different from either of the local frames of reference. Throughout this description and the accompanying claims, the terms "frame of reference", "reference frame" and "coordinate system" are used
20 interchangeably and have the same meaning.

[0032] In some embodiments of the invention, the system 20 is capable of determining the positions and/or features of objects within an additional stereo measurement field, wherein the fields of view of
25 imaging devices belonging to two different stereo measurement units overlap. In Figure 2, region 72 comprises a portion of an additional stereo measurement field within which the fields of view of imaging

20092879 63000

- 12 -

devices **52B** (which belongs to stereo measurement unit **50**) and **62B** (which belongs to stereo measurement unit **60**) overlap. Line **69** is a boundary of a total stereo measurement field which can be achieved in such an embodiment of the invention. The total system stereo measurement field for these embodiments includes the union of stereo measurement fields **58**, **68** and additional stereo measurement region **72**.

[0033] In some cases, the accuracy with which the position and/or features of an object can be determined may be improved by determining the position of the object using data which includes data from sensors on different stereo measurement units. Such measurements can benefit from a longer baseline B between stereo measurement units (50 and 60).

[0034] In some embodiments of the present invention, stereo measurement units (50 and 60) are relatively far apart (i.e. baseline B is relatively large) in comparison to the baselines b_{50} and b_{60} of stereo measurement units (50 and 60). For example B may be more than $2 \frac{1}{2}$ times greater than either of b_{50} and b_{60} .

20 [0035] For the sake of simplicity and ease of description, three-dimensional imaging system 20 is depicted in Figure 2 as having only two stereo measurement units (50 and 60). In general, multi-nodal, three-dimensional imaging system 20 requires at least two measurement units, but may include more than two depending on accuracy
25 requirements, coverage requirements and financial constraints of a particular three-dimensional imaging application.

- 13 -

[0036] System 20' of Figure 8 comprises three stereo measurement units (40, 50 and 60). Each stereo measurement unit (40, 50 and 60) has a stereo measurement field (48, 58 and 68), within which the positions and/or features of objects can be determined from data obtained by the sensors of the stereo measurement unit. Each of the stereo measurement fields has a region of overlap with the stereo measurement field of at least one other stereo measurement unit. For example, stereo measurement fields 48 and 58 overlap within region 70₄₈₋₅₈. Stereo measurement fields 58 and 68 overlap within region 70₅₈₋₆₈.

10

[0037] In some embodiments of the invention, system 20' is adapted so that it can be calibrated to permit determination of the locations and/or features of objects, which are located within regions within which the fields of view of imaging devices belonging to different stereo measurement units overlap (even if the object is not within the stereo measurement field of any one of the stereo measurement units). For example, in such embodiments, the positions and/or features of object in region 72₄₈₋₅₈ may be determined from data from imaging device 42B of stereo measurement unit 40 and imaging device 52A of stereo measurement unit 50.

15
20

[0038] For clarity of explanation, each stereo measurement unit (for example, unit 50) is depicted in Figure 2 as having only two imaging devices 52A and 52B. In general, each stereo measurement unit 50 has two or more imaging devices 52A and 52B. For example, each stereo measurement unit may comprise a DIGICLOPS™ stereo vision camera unit available from Point Grey Research Inc. of Vancouver, Canada.

25

2009027-6282600T

- 14 -

Such stereo vision camera units each comprise three digital cameras mounted rigidly to one another in a triangular configuration as shown in Figure 7. Stereo camera measurement unit 50' of Figure 7 includes a top camera 52A, left camera 52B and right camera 52C. Each camera
5 comprises a CCD sensor (not shown) which generates a digital image of objects within the camera's field of view.

[0039] As discussed above, the invention can be practised using measurement units in which the imaging devices comprise sensors other
10 than cameras. In some alternative embodiments of the invention, three-dimensional imaging system 20 or individual stereo measurement units (50 and 60) may even comprise imaging devices of different types. For example, in a stereo measurement unit, one imaging device may
15 comprise a digital camera, while another imaging device comprises an acoustic sensor.

[0040] Figure 3 is a high level schematic representation of a method for processing the image data 80 obtained from the individual
20 imaging devices (52A, 52B, 62A and 62B) of system 20 of Figure 2.

[0041] Image device data 80 may be acquired simultaneously from all imaging devices (52A, 52B, 62A and 62B) or may be acquired sequentially from each such device. After acquisition and digitization (if
25 required), image device data 80 is directed to a processor (not shown in Figure 3). The image device data 80 may be carried over any suitable interface. The IEEE 1394 and USB interfaces are but two examples. Within the processor, a calibration process 30 processes image device

20080501 152825001

- 15 -

data 80 to determine calibration function(s) for determining the positions and/or features of objects in a common frame of reference. The calibration function(s) depend upon the spatial positions and orientations of stereo measurement units 50 and 60 relative to one another.

- 5 Calibration process 30 produces a set of calibration parameters 81 which characterize the calibration function(s).

[0042] Depending on the application requirements, calibration process 30 may be performed continuously, each time that system 20 starts up, at preselected time intervals, on the occurrence of a preselected event or when specified by external commands. Data integration process 32 uses calibration parameters 81 along with image device data 80, to determine the positions and/or features of objects in the common coordinate system. Some of the objects or features may be only within the stereo measurement field of stereo measurement unit 50, others may be only within the stereo measurement field of stereo measurement unit 60. In some embodiments of the invention other objects may be within additional stereo measurement field 72, but not within the stereo measurement field of either one of the stereo measurement units.

20

[0043] In data integration process 32, a single "point cloud" which specifies the positions and/or features of objects is obtained from image data 80 and calibration parameters 81. A point cloud is a set of points having positions specified in a common coordinate system.

25

[0044] Output process 34 receives data specifying a single three-dimensional point cloud with a common coordinate system. Output

20090802 10092279-030802

- 16 -

process 34 may then provide three-dimensional data relating to the objects in the overall system stereo measurement field for any desired application. Three-dimensional data may include characteristics of objects (e.g. size, position and orientation) that are in the overall system stereo measurement field and output process 34 may, for example, involve tracking such object characteristics over time.

[0045] The method of Figure 3 may be performed by a data processor executing software instructions, by hardware configured to perform the method or by a combination of hardware and software. The method may be performed at a central location, such as within a computer 90 (see Figure 2) or may be distributed. Calibration process 30, data integration process 32 and output process 34 may be performed in a single hardware device and their functionality may be achieved by a single software program. It will be appreciated by those skilled in the art, however, that these processes may be performed by distinct hardware or software elements in communication with one another and that these distinct hardware or software elements may be located within individual stereo measurement units or in a separate computer system 90.

[0046] In one embodiment of the invention, calibration process 30, data integration process 32 and output process 34 are implemented using a conventional computer 90 equipped with suitable software and interface hardware (see Figure 2). In this manner computer 90 exists as an independent entity which can be linked to the rest of three-dimensional imaging system 20 as and when required. Alternatively, calibration process 30, data integration process 32 and output process 34

- 17 -

may be implemented using an embedded microprocessor system (not shown) that is an integral part of one of the elements of three-dimensional imaging system 20. An embedded system implementation can be well adapted to applications involving continual monitoring or tracking or in environments where the positions or alignments of stereo measurement units 50 and 60 are prone to change.

[0047] In a further alternative embodiment, each of stereo measurement units 50 and 60 is equipped with its own on-board hardware or hardware-software combination configured to perform calibration process 30. In such embodiments of the invention, data integration process 32 and output process 34 may be implemented using a central computer 90 or an embedded processor (not shown). In such embodiments of the invention, stereo measurement units 50 and 60 may output raw image data or may process the image data and output processed data. For example, stereo measurement units may output data specifying the locations of points in a point cloud in some coordinate system.

[0048] In yet another embodiment, a portable calibration device (not shown) containing a processor for performing calibration process 30 and suitable I/O hardware is used. The portable calibration device may be connected to stereo measurement units in turn and used to sequentially calibrate all of the stereo measurement units 50 and 60. After calibration process 30 is performed by such a detachable device, data integration process 32 and output process 34 may be implemented on a central computer 90 or an embedded processor which receive calibration

- 18 -

information from the portable calibration device. A portable calibration device may comprise a portable computer processor, programable hardware, such as a field programmable gate array or the like.

- 5 [0049] Figure 4 is a flow chart depicting a method for calibrating a three-dimensional imaging system 20 according to one embodiment of the invention. The result of calibration process 30 is calibration parameters 81, which comprise information specifying a set of one or more transformations. The set of transformations transforms positions and/or features of objects derived from the image data 80 from each stereo measurement unit (50 and 60) to a common frame of reference. The transformations depend upon the relative positions and orientations of each of the stereo measurement units (50 and 60) with respect to the common frame of reference. The positions and orientations of each of the stereo measurement units may vary in the common frame of reference in six degrees of freedom (three translational and three orientational).

- [0050] Referring to Figures 2 and 4, an internal calibration process 95 performs internal calibration of stereo measurement units 50 and 60.
- 20 Internal calibration process 95 obtains calibration information as between individual imaging devices (52A, 52B, 62A and 62B) within each such stereo measurement unit 50 and 60. This may be done according to any suitable technique. For example, calibrating a stereo measurement unit may involve placing a known object at a known location relative to the stereo measurement unit, obtaining images of the object from a plurality of imaging devices of the stereo measurement unit, locating an image of a reference point on the object in each of a
- 25

- 19 -

plurality of the images and deriving calibration information from the known actual location of the reference point and the locations of the image of the reference point. Internal calibration process 95 may comprise retrieving previously obtained internal calibration information.

5

[0051] When internal calibration information has been obtained from all of the stereo measurement units in system 20, the internal calibration information can typically be retained for a relatively long period of time without recalibration, since the imaging devices of stereo measurement units 50 and 60 are typically rigidly and stably mounted to one another.

[0052] After internal calibration process 95, the function (transformation) which permits the position of a reference point to be determined in the local coordinate system of stereo measurement unit 50 from image data obtained by imaging devices 52A and 52B is known. Similarly, the function (transformation) which permits the position of a reference point to be determined in the local coordinate system of stereo measurement unit 60 from image data obtained by imaging devices 62A and 62B is known. These functions (transformations) are equivalent to function F of equation (1). Thus, after the completion of process 95, the processor can determine the three-dimensional position and/or features of an object within the stereo measurement field (58 or 68) of the corresponding stereo measurement unit (50 or 60) relative to a reference frame local to that stereo measurement unit.

25

20080601 64825001

- 20 -

[0053] Image acquisition process 100 involves obtaining images of one or more calibration targets and providing resulting image device data 80 obtained by each of imaging devices 52A, 52B, 62A and 62B to a processor. One or more calibration targets is placed in overlapping stereo measurement field 70. Each calibration target is a known object having one or more features. The calibration target is within stereo measurement fields 58 and 68 of both stereo measurement units 50 and 60.

[0054] Image data 80 may be transferred to the processor simultaneously for all stereo measurement systems 50 and 60 or at different times for different stereo measurements systems. The image data 80 may be acquired simultaneously or image device data 80 from different stereo measurement units (50 and 60) may be obtained at different times. The choice of whether to acquire all image device data 80 simultaneously or sequentially and whether to deliver the image device data 80 to a processor simultaneously or at different times depends on a number of factors, including:

- the processing power of the processor performing calibration process 30, data integration process 32 and output process 34;
- the bandwidth of the interface between imaging devices and the processor;
- the resolutions of imaging devices 52A, 52B, 62A and 62B, (this determines the total amount of data that must be processed); and,
- whether the stereo measurement units 50 and 60 or calibration targets employed are likely to move.

If the processor speed is relatively low and the resolution of imaging devices 52A, 52B, 62A and 62B is relatively high, then it can be

- 21 -

desirable to acquire image device data 80 sequentially. In contrast, if the calibration target(s) and/or stereo measurement units 50 and 60 might move during calibration, then image device data 80 should be acquired simultaneously to maintain accurate correspondence between the image
5 device data 80 captured from imaging devices 52A, 52B, 62A and 62B.

[0055] Feature localization and extraction process 110 involves locating relevant calibration targets and particular calibration features from image device data 80. During feature localization and extraction
10 process 110, image device data 80 from the field of view of each imaging device 52A, 52B, 62A and 62B is processed to locate at least one calibration feature of a calibration target. Assuming that a calibration target having a calibration feature is within stereo measurement field 58, the three-dimensional location of that calibration feature may be
15 determined in the local coordinate system of stereo measurement unit 50 from image data 80 of imaging devices 52A and 52B and the internal calibration information obtained in internal calibration process 95. Similarly if the same calibration feature is also in stereo measurement field 68, the three-dimensional location of the calibration feature may
20 determined relative to the local coordinate system of stereo measurement unit 60 from image data 80 of imaging devices 62A and 62B and the internal calibration information obtained in internal calibration process 95. Feature localization process 110 may be implemented in various ways.

25

[0056] A first embodiment of the invention depicted in Figure 2 involves a configuration where stereo measurement units 50 and 60 are

203060 62826001

- 22 -

- similarly oriented. In such an embodiment, a suitable calibration target may comprise a planar object. Examples of suitable calibration targets include planar surfaces marked with one or more circular dots, planar surfaces marked with one or more squares or rectangles, "checker-board patterns" and the like. Preferably, the markings on the planar surface contrast with the surface.

- [0057]** For a single dot shaped calibration target, the calibration feature extracted during feature localization and extraction process 110 may be the center of the dot. The three-dimensional position of the calibration feature (i.e. the center of the dot) may be extracted from image device data 80 in the local coordinate systems for each stereo measurement unit 50 and 60.
- [0058]** In a second embodiment shown in Figure 9, stereo measurement units 50 and 60 are oriented in substantially different directions. With such a configuration, a planar dot calibration target may produce spurious results, because the dot may appear distorted from the perspective of one or more of imaging devices 52A, 52B, 62A and 62B.
- In such cases, it may be preferable to use a calibration target 75 comprising a spherical object. The calibration target preferably has a surface which contrasts with the background. For example, the calibration target may comprise a black sphere imaged against a white background. A spherical calibration target has the same appearance when viewed from any angle. With a spherical calibration target, the calibration feature extracted during feature localization and extraction process 110 may be the center of the sphere. Spherical calibration targets

- 23 -

are also useful where imaging devices **52A**, **52B**, **62A** and/or **62B** comprise laser range-finding devices (not shown).

- [0059]** Other known objects may be used for calibration targets. For example, in the embodiment of Figure 10, where the stereo measurement units **50** and **60** are oriented to oppose one another, a suitable calibration target may comprise a two-sided planar dot. In this case the calibration feature may comprise the center of the dot.
- [0060]** In another embodiment of the feature localization and extraction process **110**, specialized calibration targets may not be required. Such an embodiment may make use of features localized and extracted from ordinary scenes. One example, of such a feature extraction technique is known as Scale Invariant Feature Transform ("SIFT"). SIFT technology is described in David G. Lowe, *Object Recognition for Local Scale-Invariant Features* (Proc. of International Conference on Computer Vision, Corfu (Sep. 1999)), which is hereby incorporated by reference. In general, any discernable feature contained in the stereo measurement field of the stereo measurement units sought to be calibrated may be used as a calibration feature using SIFT. Depending on the type of feature extracted, the information determined from the calibration feature using this type of technique may include three-dimensional position and/or orientation information.
- [0061]** Yet another embodiment of the feature localization and extraction process **110**, involves a three-dimensional model-based object recognition algorithm capable of determining the three-dimensional

- 24 -

position and the orientation of a predefined object. Such a system may employ a measurement unit comprising monoscopic or stereoscopic image sensor to obtain image data from within its field of view. The model-based-object recognition algorithm then compares that image data to data about one or more predefined objects. In an embodiment employing model-based-object recognition, data specifying one or more calibration features may permit the location and orientation of a three-dimensional calibration target to be determined. Such information can be a sufficient basis for determining the transformations in step 130 (see further discussion below).

[0062] Figure 5 is a flowchart that depicts an example method for feature localization and extraction. Feature localization and extraction may be performed separately for each stereo measurement unit. In the example of Figure 5, the calibration target comprises a planar dot and the center of the dot is the calibration feature.

[0063] In feature location step 112, the raw image data 80 from a reference one of the imaging devices (for example, imaging device 52A) is received and the dot is located within that image data. Step 112 outputs the two-dimensional location of the dot center in the image produced by the reference imaging device. The dot may be located by way of any of a variety of feature recognition methods. For example, a Hough transform may be performed on the image data 80 to locate circles (i.e. dots) and their centers represented in the image data. Hough transforms are described in U.S. Patent No. 3,069,654, which is hereby incorporated by reference. After the dot is located, the two-dimensional location of its

- 25 -

center (i.e. the calibration feature) can be refined by employing the same Hough transform to search within a restricted search area.

[0064] Optional estimation step 114 employs knowledge of the size
5 of the dot located in step 112, the internal calibration parameters of stereo measurement unit 50 and the focal lengths of the various individual imaging devices (52A and 52B) within that stereo measurement unit 50 to estimate the two-dimensional position at which
the dot is likely to be found within the image data 80 of the other
10 imaging devices in the stereo measurement unit (in the illustrated embodiment, the only other imaging device is imaging device 52B).

[0065] If estimation step 114 is not performed then the two-
dimensional location of the calibration feature may be located within the
15 image data 80 of the non-reference imaging device 52B in a different way, for example by using a Hough transform process similar to step 112. Estimation step 114 is useful to reduce the amount of processing required to locate the calibration feature in two-dimensions within the image data 80 of the non-reference imaging device 52B. Estimation step
20 114 particularly increases efficiency in embodiments (not shown) where there are more than two imaging devices in a stereo measurement unit.

[0066] Optionally the two-dimensional location of the dot center
within the image data 80 of each individual imaging device (52A and
25 52B) may be further refined in step 116 to sub-pixel accuracy. Step 116 may comprise obtaining a histogram of binarized pixel values in the

- 26 -

region of the dot and then determining a best fit of the binarized pixels to a circle.

[0067] Triangulation process 118 involves using the locations of the two-dimensional calibration features from each imaging device 52A and 52B to determine the three-dimensional location of the calibration feature. This may be done using conventional triangulation techniques. For example, the following stereo triangulation equations may be used to locate the center of a calibration target:

10
$$z = \frac{b \cdot f}{d} \quad (2)$$

$$x = \frac{z \cdot u}{f} \quad (3)$$

$$y = \frac{z \cdot v}{f} \quad (4)$$

where:

- (x,y,z) is the location of the dot center in the three-dimensional local coordinate system of the stereo measurement unit;
- 15 • (u,v) are the coordinates of the dot center in the reference camera;
- b is the baseline;
- d is the disparity value, reflecting the difference in two-dimensional position of the dot center in the non-reference camera as compared to its two-dimensional position in the reference camera; and,
- 20

- 27 -

- f is the focal length of the cameras (in this example the cameras are each assumed to have the same focal length. This is convenient but not absolutely necessary).

5 [0068] The method of Figure 5 is repeated for each stereo measurement unit involved in the calibration.

[0069] Returning to Figure 4, step 120 determines whether enough information has been obtained to determine a transformation which
10 permits position information obtained from all of two or more stereo measurement units to be expressed in a common coordinate system. If not, then feature localization and extraction loop 125 is repeated for different calibration features until sufficient information has been obtained.

15 [0070] In general, there are six degrees of freedom (i.e. three degrees of translational freedom and three degrees of orientational freedom) which specify the position and orientation of the local reference coordinate systems of two stereo measurement units relative to one
20 another. The number of repetitions of image acquisition process 100 and feature localization and extraction process 110 depend on the type of calibration targets and calibration features used. For example, where a calibration feature is a point in space (i.e. the center of a dot), the extraction of one such calibration feature in process 110 yields
25 information about only three translational degrees of freedom. In general it is necessary to know the locations of three points in space (or to have equivalent information from which those locations could be determined)

20030000 " 628260001

- 28 -

in each of two local coordinate systems to be able to define a transformation between the coordinate systems (or a set of transformations to a common coordinate system). The three points may comprise different reference points on a single calibration target (e.g. a corner and points on edges of a cube), reference points on a plurality of different calibration targets, or the same reference point of a single calibration target which is moved to at least three positions within region 70.

10 [0071] In some applications, the number of calibration features extracted may exceed the minimum number of calibration features required. For example, where the calibration feature is a point, the minimum number of extracted calibration features required to obtain the six degrees of freedom is three. However the accuracy with which the
15 transformations to a common frame of reference may be determined may be improved by using additional calibration features. The additional calibration features are redundant in that they can be omitted and a set of transformations to a common frame of reference could still be determined. For this type of application, the threshold of query 120 can
20 be set arbitrarily high.

[0072] Step 130 uses the information obtained in loop 125 to generate a set of transformations from the local coordinate systems of the plurality of stereo measurement units to a common coordinate system.
25 The common coordinate system may be the same as the local coordinate system for one of the stereo measurement units. Calibration parameters 81 may comprise information regarding the positions and orientation of

20080501 15:43:25.000

- 29 -

stereo measurement units **50** and **60** along with the set of transformations calculated in step **130**.

[0073] The transformation(s) produced by step **130** may, for example, be expressed as matrices, which transform position information in the local coordinate systems of stereo measurement units **50** and **60** to position information in the common coordinate system.

[0074] One way to obtain the transformation to a common coordinate system is to perform an iterative least-squares algorithm. Figure 11 depicts one possible method which may be used to obtain a transformation in step **130**. In the embodiment of Figure 11, one stereo measurement unit (for example, unit **50**) is selected to be the "master" unit. The common coordinate system is chosen to be the same as the local coordinate system for the master unit. Transformations are calculated for one or more "slave" stereo measurement units (for example, unit **60**). The calculated transformation(s) transform the reference coordinate system of the slave unit **60** to that of the master unit **50**. In three-dimensional imaging systems employing three or more stereo measurement units (see Figure 8), a process similar to that of Figure 11 may be performed for each slave stereo measurement unit.

[0075] In a further alternative embodiment, there need be no master and slave stereo measurement units and a transformation could be applied to image data from every stereo measurement unit.

- 30 -

[0076] For the purposes of explaining the transformation calculation process of Figure 11, it is assumed that stereo measurement unit 50 is the master unit and that a transformation is sought for the image data 80 from stereo measurement unit 60. The calculation and optimization process depicted in Figure 11 is an iterative process that begins by obtaining an estimated transformation in step 132. The estimated transformation may be based upon on system parameters. In the alternative, the estimated transformation may be some predetermined transformation.

10

[0077] A master calibration point set specifying the locations of a plurality of calibration features in the local coordinate system of the master stereo measurement unit is obtained in feature localization and extraction process 110 as described above. Similarly, a slave calibration point set specifying the locations of the plurality of calibration features in the local coordinate system of the slave stereo measurement units is obtained for each slave stereo measurement unit as described above.

15

[0078] For each slave measurement unit, the current transformation (which, in the initial iteration is the estimated transformation) is applied to the corresponding slave calibration point set to yield a transformed slave point set in step 137. In error calculation step 133, the transformed slave point set is compared to the master calibration point set and an error in the current transformation is determined. In some embodiments the error is expressed as an error matrix. The error may be determined by evaluating the distance between each transformed slave point and a corresponding calibration point in the master point set. A residual error,

20

25

20080802 12:52 FAX 604 681 4081

- 31 -

such as the sum of the squared distance errors for each calibration point, indicates the "quality of fit" of the current transformation. In each iteration a record is kept of the residual error.

5 [0079] The error is used to generate a correction, which is applied to the current transformation in step 135. The correction may be, for example, represented as a correction matrix. The correction matrix may be expressed as a function of the error matrix and the matrix of partial derivatives of the current transformation as applied to each calibration
10 point in the slave point set. The correction can be determined in a computer or other data processing mechanism using well-known mathematical techniques, such as an iterative least-squares algorithm.

[0080] Once the error is determined in error calculation step 133,
15 the correction matrix is used to generate a new current transformation matrix in transformation update step 135. This may involve multiplying a current transformation matrix by a correction matrix. Transformation application steps 137 and error calculation step 133 are repeated. Step 134 determines whether the residual error calculated in step 133 is
20 unchanged or has increased compared to the previous iteration. If the residual error is unchanged or has increased then the method proceeds to step 139.

[0081] The process then determines whether there are any
25 redundant calibration features for which the positional error between the calibration feature in the master calibration point set and the corresponding feature in the transformed slave point set is statistically

20090906 09:03:00

- 32 -

outlying. This may be done, for example, by computing an average and standard deviation of the positional error over the calibration features in the two point sets. Calibration features for which the positional error is more than, for example, two standard deviations away from the average
5 may be considered to be statistically outlying. If there are statistically outlying calibration features as determined in step 138, then they are removed from consideration in step 136 and then the transformation is returned to transformation application step 137 for possible optimization. The transformation calculation and optimization are complete when there
10 are no redundant statistically outlying calibration features and the answer to query 138 is negative.

[0082] In a preferred embodiment, a plurality of extra calibration features are localized and extracted during feature extraction loop 125
15 (see Figure 4). If extra calibration features are localized and extracted during the feature extraction loop 125, they may be used for optimizing the calculation of the various transformations in transformation calculation step 130. Optimization techniques may include selective weighting of extracted calibration features that are located in regions of
20 interest over extracted features that are located elsewhere, to provide the calculated transformations with more accuracy in the regions of interest. Alternatively, where a large number of calibration features are extracted, optimization techniques, such as weighting or clustering, may be used to ensure that the resultant transformations do not unevenly favour a region
25 with a large number of calibration features. In addition, having extra calibration features allows for features to be dropped in step 136 (see Figure 11), if they are statistically outlying.

- 33 -

- [0083] The algorithm of calibration process 30 depicted in Figure 4 involves a single feature extraction loop 125, where all of the features are extracted prior to calculation of the transformation matrices in step 130.
- 5 Another embodiment of calibration process 30, which is depicted in Figure 6, involves continuously updating the transformation matrices in a calibration loop 145 and continually updating the calibration parameters 81. The embodiment of Figure 6 involves the same initial steps: image acquisition process 100, feature localization and extraction process 110
- 10 and feature extraction loop 125. Once enough calibration features are extracted, the system performs an original transformation calculation in step 130 and outputs calibration parameters 81 in step 140, before looping back to acquire more image data in calibration loop 145. After a first transformation calculation is made in step 130, subsequent
- 15 transformation calculations made in step 132 are used to update calibration parameters 81. In this manner, the transformation matrices and resultant calibration parameters 81 are optimized over time, by making use of a growing number of calibration features. After each transformation calculation, the updated calibration parameters 81 are
- 20 output in step 140. Calibration loop 145 may be stopped at any time, if it is determined that calibration parameters 81 are sufficiently accurate or that a sufficient number of calibration features has been extracted.

- [0084] The alternative methods depicted in Figures 4 and 6 have
- 25 various advantages in different applications. One advantage of the method depicted in Figure 6 is that it allows calibration to become more accurate over time, because a large number of calibration features can be

20090809 090000

- 34 -

- extracted. In addition, the embodiment of Figure 6 involves continuously updating calibration, such that calibration may be restored even if stereo measurement units 50 and 60 are moved. In contrast, the method depicted in Figure 4 allows a rapid calibration determination, because all of the calibration features are extracted prior to the calculation of the transformations in step 130. Because of the increased speed of calibration feature extraction, the method of Figure 4 may be preferred in situations where processing resources are at a premium.
- 10 [0085] The embodiments of calibration process 30 described above depend on locating and extracting calibration features that are contained in overlapping measurement field 70, which comprises parts of the stereo vision fields 58 and 68 (see Figure 2). An alternative method of calibration illustrated in Figure 13 uses a calibration target 100 that spans stereo measurement fields 58 and 68. Such a calibration target may comprise, for example, a pair of objects 102,104 rigidly connected by a rod 106 or other member. The calibration target has a first set of calibration features in the first stereo measurement field 58 that have a known spatial relation to a second set of calibration features in the second stereo measurement field 68. Because the calibration features in each stereo measurement field have a known spatial relationship, if locations of the first set of calibration features are known in the local coordinate system of a first measurement unit then the locations of the second set of calibration features in the same local coordinate system can be determined from the known relationship between the two sets of calibration features. The locations of the second set of calibration features in the local coordinate system of a second measurement unit (as
- 25

20030802 12:54 FAX 604 681 4081

- 35 -

determined by data acquired by the second measurement unit) can then be used to determine a set of one or more transformations to a common coordinate system. With such a calibration technique, the individual measurement fields 58 and 68 may be non-overlapping as shown in

5 Figure 13.

[0086] Another alternative calibration technique for calibrating first and second stereo measurement units which have non-overlapping stereo measurement fields involves using one or more additional stereo
10 measurement units which have stereo measurement fields which overlap with those of the first and second stereo measurement units. The additional stereo measurement(s) may be removed after calibration is completed. Assuming a configuration similar to that of Figure 2, this embodiment does not require any overlap of stereo measurement fields
15 58 and 68. The extra stereo measurement unit may be inserted into the environment in such a manner that its stereo measurement field overlaps with measurement fields 58 and 68. Calibration may then be performed in the manner described above. After calibration, the extra measurement unit may be removed. Despite the removal of the extra measurement unit
20 and the lack of overlap between measurement fields 58 and 68, measurement units 50 and 60 remain calibrated to one another as long as measurement units 50 and 60 do not move. Measurement units 50 and 60 may be used to obtain three-dimensional data in the common coordinate system.

25

[0087] Referring back to Figure 3, after calibration process 30 produces calibration parameters 81, system 20 may be used to determine

20090901 10090901

- 36 -

positions and/or features of objects within the overall system stereo measurement field. Data integration process 32 involves the application of the transformations to image data 80, in order to bring the image data 80 for stereo measurement units 50 and 60 into a common coordinate system. In data integration process 32, a single point cloud is obtained from the calibration parameters 81 and the image data 80. Output process 34 provides three-dimensional data relating to the objects in the overall system stereo measurement field for any desired application.

10 [0088] Figure 12 represents an alternative embodiment in which raw image data 80 from stereo measurement units 50 and 60 is directed to a calibration and triangulation process 30. Calibration and triangulation process 30 has a dual function involving the acquisition of calibration parameters 81 and the calculation of localized three-dimensional data 82. Calibration and triangulation process 30 may obtain calibration parameters 81 from the image data 80 according to one of the processes described above (see Figures 4 and 6). Calibration and triangulation process 30 triangulates the raw image data 80 to obtain localized three-dimensional data 82, which contains a three-dimensional point cloud of image data for each stereo measurement unit 50 and 60. Localized three-dimensional data 82 is three-dimensional image data localized to the frame of reference of the individual stereo measurement units 50 and 60. After calibration and triangulation process 30, calibration parameters 81 and localized three-dimensional data 82 are directed to data integration process 32, which uses calibration parameters 81 to transform localized three-dimensional data 82 into a single common coordinate system. Data integration process 32 of Figure 12

- 37 -

does not need to work with raw image data 80. The output process 34 of the Figure 12 embodiment may be substantially similar to that of the Figure 3 embodiment.

5 [0089] In a further alternative embodiment (not shown), triangulation of raw image data 80 to obtain localized three-dimensional data 82 may be performed in a separate triangulation process that precedes calibration process 30. In this manner, calibration is performed using only localized three-dimensional data 82.

10

[0090] The embodiments of Figures 3 and 12 have various advantages in different applications. One advantage of the method depicted in Figure 12 is that the calibration and triangulation process 30 may be conveniently performed in a data processing system which is specifically associated with or embedded into a stereo measurement unit (i.e. 50 or 60), such that stereo measurement units need only output localized three-dimensional data 82 as opposed to raw image data 80. Stereo measurement units can be made as modular devices. The method of Figure 3 is advantageous, because data integration process 34 receives raw image data 80 from individual imaging devices and may perform an algorithm for selecting the best pairing of individual imaging devices to provide most accurate three-dimensional data. For example (see Figure 2), to measure the positions and/or features of an object in stereo measurement field 68 and close to stereo measurement unit 60, imaging devices 62A and 62B may be selected; for objects in overlapping stereo measurement field 70, imaging devices 52A and 62B may be selected;

15
20
25

20030801 16:22:50

- 38 -

and for objects in region **72**, imaging devices **52B** and **62A** may be selected.

[0091] The system may optionally derive calibration information
5 for pairs of imaging devices which do not lie within the same stereo
measurement unit. This may be done by locating a number of object
features which are both within a stereo measurement field of such a pair
of imaging devices and within a stereo measurement field of a stereo
measurement unit. From the known locations of the object features (as
10 determined by the stereo measurement unit) a transformation can be
derived which provides the position of an object feature in the common
coordinate system from the positions of the object feature in image data
from the pair of imaging devices.

15 [0092] In some cases, three-dimensional measurements of the
positions of features on distant objects made by system **20** may have
increased accuracy because the separation between stereo measurement
units **50** and **60** (i.e. the baseline **B**) may be made arbitrarily large. Since
three-dimensional measurements can also be made by a single stereo
20 measurement unit, measurements can be made on objects that are
relatively close to a stereo measurement unit.

[0093] Figure 14 shows a method **200** according the invention.
Method **200** begins by providing at least two measurement units each
25 capable of locating objects within a corresponding measurement field
(block **210**). Method **200** continues by obtaining information about a
first set of calibration features (which may comprise points, lines, etc.) in

- 39 -

2080E0" 6ZB2600T

a first reference frame local to a first measurement unit (block 212). The first set of calibration features is located in a measurement field of the first measurement unit. The information obtained in block 212 is used in block 214 to determine locations of a second set of calibration features
5 relative to the first reference frame. The second set of calibration features are located in the measurement field of a second measurement unit. In block 216 locations of the second set of calibration features relative to the second reference frame is determined. In block 218 the locations of the second set of calibration features relative to the first and
10 second reference frames are used to establish a set of one or more transformations from the local reference frames to a common reference frame. In block 220 the transformations are used to express the locations of reference points (which may be in either measurement field) relative to the common reference frame.

15 [0094] In the foregoing description, image processing steps which are common and well known to those skilled in the art have been omitted for clarity. For example, the method of the invention may involve processing raw data to correct for image distortions introduced by camera
20 lenses. Such image processing techniques are known to those skilled in the field of image processing. Computer software which implements such techniques is available commercially from Point Grey Research Inc., of Vancouver, Canada.

25 [0095] Certain implementations of the invention comprise computer processors which execute software instructions which cause the processors to perform a method of the invention. The invention may also

- 40 -

be provided in the form of a program product. The program product may comprise any medium which carries a set of computer-readable signals comprising instructions which, when executed by a data processor, cause the data processor to execute a method of the invention. The program
5 product may be in any of a wide variety of forms. The program product may comprise, for example, physical media such as magnetic data storage media including floppy diskettes, hard disk drives, optical data storage media including CD ROMs, DVDs, electronic data storage media including ROMs, flash RAM, or the like or transmission-type media such
10 as digital or analog communication links.

[0096] Where a component (e.g. a software module, processor, assembly, device, circuit, etc.) is referred to above, unless otherwise indicated, reference to that component (including a reference to a
15 "means") should be interpreted as including as equivalents of that component any component which performs the function of the described component (i.e., that is functionally equivalent), including components which are not structurally equivalent to the disclosed structure which performs the function in the illustrated exemplary embodiments of the
20 invention.

[0097] As will be apparent to those skilled in the art in the light of the foregoing disclosure, many alterations and modifications are possible in the practice of this invention without departing from the spirit or scope
25 thereof. For example:

- stereo measurement units may comprise any device capable of locating objects in space including: monocular vision systems

20080601 09:22:00

- 41 -

capable of determining a location of the monocular vision system in three dimensions by determining a pose and location of a known calibration object relative to the monocular vision system, laser rangefinders, stereo cameras, and the like;

- 5 • Processor 14 or 90 may comprise a programmed computer processor executing software instructions, a hardware processor or a combination of hardware and software;
- it is a matter of design convenience whether, after the system is calibrated, the positions of points are first determined in local
- 10 coordinate systems of individual stereo measurement units to produce localized three-dimensional data 82 and then converted in data integration process 32 to a common coordinate system (i.e. the embodiment of Figure 12) or whether the positions of the points in the common coordinate system are determined in data
- 15 integration process 32 directly from raw image data 80 from individual imaging devices (i.e. embodiment of Figure 3).

Accordingly, the scope of the invention is to be construed in accordance with the substance defined by the following claims.

20090926-0000T